

MEBoost: Variable Selection in the Presence of Measurement Error

Ben Brown*

Timothy Weaver[†]

Julian Wolfson[‡]

Abstract

We present a novel method for variable selection in regression models when covariates are measured with error. The iterative algorithm we propose, Measurement Error Boost (MEBoost), follows a path defined by estimating equations that correct for covariate measurement error. Via simulation, we evaluated our method and compare its performance to the recently-proposed Convex Conditioned Lasso (Co-CoLasso) (Datta and Zou, 2017) and to the “naive” Lasso (Tibshirani, 1996) which does not correct for measurement error. Increasing the degree of measurement error increased prediction error and decreased the probability of accurate covariate selection, but this loss of accuracy was least pronounced when using MEBoost. We illustrate the use of MEBoost in practice by analyzing data from the Box Lunch Study, a clinical trial in nutrition where several variables are based on self-report and hence measured with error.

Key words: Boosting; High-Dimensional Data; Lasso; Machine Learning; Measurement Error; Variable Selection

*Division of Biostatistics, School of Public Health, University of Minnesota, MMC 303, 420 Delaware Street SE, Minneapolis, MN 55455

[†]Minneapolis Medical Research Foundation, 914 S 8th St, Minneapolis, MN 55404

[‡]Division of Biostatistics, School of Public Health, University of Minnesota, MMC 303, 420 Delaware Street SE, Minneapolis, MN 55455

1 Introduction

Variable selection is a well-studied problem in situations where covariates are measured without error. However, it is common for covariate measurements to be error-prone or subject to random variation around some mean value. Consider, for instance, a study wherein subjects report their daily food intake on the basis of a dietary recall questionnaire. There is variation from day to day in an individual’s calorie consumption, but it is also well established in the nutrition literature that there is error associated with the recall or measurement of the number of calories in a meal (Spiegelman et al., 1997; Fraser and Stram, 2012). In the usual regression setting, ignoring measurement error leads to biased coefficient estimation (Rosner et al., 1992), and hence the presence of measurement error has the potential to affect the performance of variable selection procedures.

There has been relatively little research done about variable selection in the presence of measurement error. Sørensen et al. (2012) introduced a variation of the Lasso that allows for Normal, i.i.d., additive covariate measurement error. Datta and Zou (2017) proposed the convex conditioned Lasso (CoCoLasso) which corrects for both additive and multiplicative measurement error in the normal case. Both of these methods are applicable to linear models for continuous outcomes, but do not easily extend to regression models for other outcome types (e.g., binary or count data). Meanwhile, there is a sizable statistical literature on methods for performing estimation and inference for low-dimensional regression parameters in the presence of measurement error (Rosner et al., 1992; Stefanski and Carroll, 1985; Fuller, 1987), but these approaches do not address the variable selection problem and cannot be applied in large p , small n problems.

We propose a novel method for variable selection in the presence of measurement error, MEBoost, which leverages estimating equations that have been proposed for low-

dimensional estimation and inference in this setting. MEBoost is a computationally efficient path-following algorithm that moves iteratively in directions defined by these estimating equations, only requiring the calculation (not the solution) of an estimating equation at each step. As a result, it is much faster than alternative approaches involving, e.g., a matrix projection calculation at each step. MEBoost is also flexible: the version that we describe is based on estimating equations proposed by Nakamura (Nakamura, 1990), which apply to any generalized linear model, and the underlying MEBoost algorithm can easily incorporate measurement error-corrected estimating equations for other regression models. We conducted a simulation study to compare MEBoost to the Convex Conditioned Lasso (CoCoLasso) proposed by Datta and Zou (2017) and the “naive” Lasso which ignores measurement error. We also applied MEBoost to data from the Box Lunch Study, a clinical trial in nutrition where caloric intake across a number of food categories was based on self-report and hence measured with error.

2 Background

2.1 Regression in the Presence of Covariate Measurement Error

Our discussion of measurement error models draws heavily from Fuller (1987). When modeling error the covariates can be treated as random or fixed values. Structural models consider the covariates to be random quantities and functional models consider the covariates to be fixed (Buonaccorsi, 2010). We consider a structural model. Let $\mathbf{Y} = \mathbf{X}\beta + \epsilon$, where \mathbf{X} is a (random) matrix of covariates of dimension $n \times p$, β a vector of coefficients of length p , ϵ is a vector of Normally distributed i.i.d. random errors of length n , and \mathbf{Y} is the resultant outcome vector also of length n . In an additive measurement error model,

we assume that what is observed is not \mathbf{X} but rather the “contaminated” or “error-prone” matrix $\mathbf{W} = \mathbf{X} + \mathbf{U}$ where \mathbf{U} a random $n \times p$ matrix.

When a model is fit that ignores measurement error, i.e. it assumes that the true model is $\mathbf{Y} = \mathbf{W}\beta_{\mathbf{W}} + \epsilon$, the resulting estimates $\hat{\beta}_{\mathbf{W}}$ are said to be *naive* and satisfy

$$E[\hat{\beta}'_{\mathbf{W}}] = \beta'(\Sigma_{\mathbf{X}\mathbf{X}} + \Delta)^{-1}\Sigma_{\mathbf{X}\mathbf{X}} \quad (1)$$

where β is the true coefficient vector, $\Sigma_{\mathbf{X}\mathbf{X}}$ is the covariance matrix of the covariates and $\Delta \equiv \Sigma_{UU}$ is the covariance matrix of the measurement error. In the case of linear regression with a single covariate, (1) simplifies to an attenuating factor that biases the coefficient estimates towards zero. However, with multiple covariates the bias may increase, decrease, and even change the sign of the estimated coefficients. Notably, measurement error affecting a single covariate can bias coefficient estimates in all of the covariates, even those that are not measured with error (Buonaccorsi, 2010).

2.2 Lasso in the Presence of Measurement Error

Sørensen et al. (2012) analyze the Lasso (Tibshirani, 1996) in the presence of measurement error by studying the properties of

$$\hat{\beta}_{Lasso, \lambda_n} = \underset{\alpha}{\operatorname{argmin}} (||\mathbf{Y} - \mathbf{W}\alpha||_2^2 + \lambda_n ||\alpha||_1) \quad (2)$$

$\hat{\beta}_{Lasso, \lambda_n}$ is asymptotically biased when $\lambda_n/n \rightarrow 0$ as $n \rightarrow \infty$ since $E[\hat{\beta}'_{Lasso, \lambda_n}] = \beta'(\Sigma_{\mathbf{X}\mathbf{X}} + \Delta)^{-1}\Sigma_{\mathbf{X}\mathbf{X}}$. Notice this is the same bias that is introduced when naive linear regression is performed on observed covariates. Sørensen et al. (2012) derive a lower bound on the magnitude of the non-zero coefficient elements below which the corresponding covariate will not be selected, and an upper bound on the L_1 estimation error $||\hat{\beta}_{\mathbf{W}} - \beta||_1$. They show that with increasing measurement error the lower bound increases, i.e., increasing measurement error adds non-informative noise to the system and so for the signal

associated with the relevant covariates to be identified the signal must increase. Increased measurement error also leads to an increase in the upper bound of the estimation error. Sign consistent selection is also impacted by the presence of covariate measurement error. Sørensen et al. (2012) set a lower bound on the probability of sign consistent selection in this setting. The result requires that the *Irrepresentability Condition with Measurement Error* (IC-ME) holds. The IC-ME requires that the measurements of the relevant and irrelevant covariates have limited correlation, relative to the size of the relevant measured covariate correlation. Note the sample correlation of the irrelevant covariates is not considered. By studying the form of the lower bound, it can be concluded that (at least when using the Lasso) measurement error introduces a greater distortion on the selection of irrelevant covariates than it does in the selection of relevant covariates.

Sørensen et al. (2012) introduced an iterative method to obtain the Regularized Corrected Lasso with constraint on the radius R :

$$\hat{\beta}_{RCL} = \arg \min_{\|\beta\|_1 < R} \{ \|y - W\beta\|_2 - n\beta' \Delta\beta + \lambda \|\beta\|_1 \} \quad (3)$$

The main results of their simulation study were consistent with their analytical results, namely that the corrected Lasso had a slightly lower selection rate for the true covariates than the naive Lasso, but was also more conservative in including irrelevant covariates. Further, the prediction error, as measured by both $\|\hat{\beta} - \beta\|_1$ and $\|\hat{\beta} - \beta\|_2$, was lower for the corrected Lasso.

The major drawback of the corrected Lasso method is that it is very computationally intensive, involving an iterative calculation where each step involves a projection of an updated $\hat{\beta}$ onto the L_1 -ball for a given radius R . The iterative process must be conducted for each fixed value of the radius R . The selected values of R provide a path of possible solutions for $\hat{\beta}_{RCL}$. Hence, the approach seems impractical for large-scale problems and

for repeated application in a simulation study.

2.3 The Convex Conditioned Lasso (CoCoLasso)

A recent paper by Datta and Zou (2017) proposes an alternative approach which they refer to as the Convex Conditioned Lasso (CoCoLasso). Consider the following reformulation of the Lasso problem,

$$\hat{\beta}_L(\lambda) = \arg \min_{\beta} \frac{1}{2} \beta' \Sigma \beta - \rho' \beta + \lambda \|\beta\|_1, \quad (4)$$

The CoCoLasso is based on the Loh and Wainwright (2012) corrections for the predictor-outcome correlation ρ and variance matrix Σ in the presence of measurement error. When error-prone covariates \mathbf{W} are measured in place of \mathbf{X} , we can get corrected estimates $\tilde{\rho}$ and $\hat{\Sigma}$:

$$\tilde{\rho} = \frac{1}{n} \mathbf{W}' \mathbf{Y} \quad \hat{\Sigma} = \mathbf{W}' \mathbf{W} - \delta^2 I \quad (5)$$

where δ^2 is the (assumed known) variance in the measured \mathbf{W} . These estimators are unbiased. A measurement error corrected Lasso estimate could then be derived by substituting $\tilde{\rho}$ and $\hat{\Sigma}$ into (4). The problem with this idea is that the corrected matrix $\hat{\Sigma}$ may not be a valid covariance matrix, since it is possible to be non positive semi-definite. If $\hat{\Sigma}$ has a negative eigenvalue, then this Lasso function would be non-convex and unbounded. To overcome this obstacle, the key to the CoCoLasso is calculating the projection of $\hat{\Sigma}$ onto the space of positive definite matrices:

$$(\hat{\Sigma})_+ = \arg \min_{\Sigma \geq 0} \|\hat{\Sigma} - \Sigma\|_{max} \quad (6)$$

The CoCoLasso then solves a standard Lasso problem in which $\hat{\Sigma}$ and ρ with the corrected values from (5) and (6), yielding the CoCoLasso estimator:

$$\hat{\beta}_C(\lambda) = \arg \min_{\beta} \frac{1}{2} \beta' (\hat{\Sigma})_+ \beta - \tilde{\rho}' \beta + \lambda \|\beta\|_1. \quad (7)$$

When $\hat{\Sigma}$ is not positive definite, the projection from (6) can be challenging to compute. However, the projection only needs to be done once, unlike the Sørensen et al. (2012) correction which requires a projection at each iteration.

3 MEBoost: Measurement Error Boosting

Our proposed variable selection algorithm, MEBoost (**M**easurement **E**rror **B**oosting), is based on an iterative functional gradient descent type algorithm that generates variable selection paths. The key idea is that, instead of following a path defined by the gradient of a loss function (e.g., the likelihood), the “descent” follows the direction defined by an estimating equation $\mathbf{g}(\mathbf{Y}, \mathbf{X}, \beta)$. The algorithmic structure of MEBoost is shared with ThrEEBoost (Thresholded Estimating Equation Boost, Brown et al. (2016)), a general-purpose technique for variable selection based on estimating equations. While ThrEEBoost described an approach to performing variable selection in the presence of correlated outcomes by leveraging the Generalized Estimating Equations (Liang and Zeger, 1986), MEBoost achieves improved variable selection performance in the presence of measurement error by following a path defined by a measurement error corrected score function due to Nakamura which is described in Section 3.1. Nakamura’s approach is applicable to linear regression models with additive or multiplicative measurement error. Closed-form corrected score functions are also derived for Poisson, Gamma, and Wald regression. By using this family of corrected score functions, the MEBoost algorithm is more broadly applicable than the corrected Lasso and CoCoLasso, neither of which is obviously generalizable beyond linear regression.

3.1 Corrected Score Function

Nakamura (1990) proposed a set of corrected score functions for performing estimation and inference in the generalized linear regression model where covariates are subject to additive measurement error with known variance matrix Δ . In general, the corrected score function S^* based on the covariates measured with error (\mathbf{W}), has the expectation equal to the score function, S , based on the true covariates (\mathbf{X}). For the normal linear model, Nakamura proposed the following correction to the negative log-likelihood to account for measurement error:

$$l^*(\mathbf{Y}, \mathbf{W}, \beta)' = -\frac{n}{2} \log(2\pi) - n \log(\sigma) - \frac{1}{2\sigma^2} \sum [(y_i - \beta' x_i)^2 - \beta' \Delta \beta] \quad (8)$$

Differentiating (8) with respect to β , we obtain the corrected score function:

$$S^*(\mathbf{Y}, \mathbf{W}, \beta)' = S(\mathbf{Y}, \mathbf{W}, \beta)' + n\sigma^{-2}\beta' \Delta. \quad (9)$$

In this case the corrected score function is the 'naive' score function, $S(\mathbf{Y}, \mathbf{W}, \beta)'$, with a measurement error correction determined by the sample size, model error, measurement errors, and the coefficient value: $n\sigma^{-2}\beta' \Delta$. The naive score function is the score function from the true model calculated with the measured covariates:

$$S(\mathbf{Y}, \mathbf{W}, \beta, \sigma) = \sigma^{-2} (\mathbf{W}' \mathbf{Y} - \mathbf{W}' \mathbf{W} \beta) \quad (10)$$

The corrected variance estimate will be calculated as $\partial l^* / \partial \sigma = 0$, which in the normal case is:

$$\hat{\sigma}^2 = n^{-1} (\mathbf{Y} - \mathbf{W} \beta^*)' (\mathbf{Y} - \mathbf{W} \beta^*) - \beta^{*'} \Delta \beta^*. \quad (11)$$

Similarly to the corrected score function, the corrected variance estimate is the naive variance estimate, $n^{-1} (\mathbf{Y} - \mathbf{W} \beta^*)' (\mathbf{Y} - \mathbf{W} \beta^*)$, with a measurement error correction. The correction reduces the estimated variance, thus subtracting the noise introduced by the measurement error. In the variance case the correction factor is determined only by the true coefficient vector and the measurement error variance.

As another example, the correction for Poisson distributed data is the following:

$$S(\mathbf{Y}, \mathbf{W}, \beta) = \sum [y_k w_k - (w_k - \Delta\beta) \exp(\beta' w_k - \beta' \Delta\beta/2)] \quad (12)$$

which we apply in our data application (see Section 5). Nakamura (1990) also provides corrections for multiplicative measurement error in linear regression, as well as measurement error in Gamma and Wald regression. In what follows, we use the normal linear additive measurement error corrected score function as part of an iterative path-following algorithm that performs variable selection in the presence of covariate measurement error.

3.2 The MEBoost Algorithm

Our proposed variable selection algorithm, MEBoost, consists of applying ThrEEBoost with the corrected score function and corrected variance estimate described in the previous section. Algorithm 1 summarizes the MEBoost procedure.

Let $\zeta \in [0, 1]$ be the fixed thresholding parameter. Starting with a β estimate of $\mathbf{0}$ and a $\hat{\sigma}^2 = 1$, the corrected score function \mathbf{S}^* is calculated at these values, and the magnitude of each component of $\nu \equiv \mathbf{S}^*$ is recorded. The indices of elements to update are identified by a thresholding rule, $J_t = \{j : |\nu_j| \geq \zeta \cdot \max_j |\nu_j|\}$. The next point in the variable selection path, $\beta^{(1)}$, is obtained by adding a small value, γ , to each of these elements in the direction corresponding to the signs of each ν_j for $j \in J_t$. This updated $\beta^{(1)}$ is used to calculate an updated corrected $\sigma^{2(1)}$. The algorithm continues for T iterations, where T is typically chosen to be large (e.g., 1,000).

The parameters γ , T and ζ interact to determine the specific variable selection path that results from the algorithm. The smaller the value of γ the smaller the distance between β estimates on the selection path, while a larger value of γ leads to larger jumps in the selection path. Ideally, a very small value of γ would be used, but if $\|\beta\|_1$ is large,

Algorithm 1 MEBoost

procedure MEBOOST

Set $\beta^{(0)} = \mathbf{0}$

Set $\sigma^{2,(t=0)} = 1$

for $t = 0, \dots, T$ **do**

 Compute $\nu = \mathbf{S}^*(\mathbf{Y}, \mathbf{W}, \beta)_{\beta=\beta^{(t-1)}}$

 Identify $J_t = \{j : |\nu_j| \geq \zeta \cdot \max_j |\nu_j|\}$

for all $j_t \in J_t$ **do**

 Update $\beta_{j_t}^{(t)} = \beta_{j_t}^{(t-1)} + \gamma \text{sign}(\nu_{j_t})$

 Set $\sigma^{2,(t)} = n^{-1} (\mathbf{Y} - \mathbf{W}\beta^{(t)})' (\mathbf{Y} - \mathbf{W}\beta^{(t)}) - \beta^{(t)'} \Delta\beta^{(t)}$

a large number of iterations, T , may be required to generate a selection path. This of course is the trade-off one is required to make when determining the step size. A selection path incremented by only a small value is preferable to a path which takes large steps, but the time required for a large number of iterations may become prohibitive. With each of the t iterations those elements of the coefficient vector that are still of size zero have not been selected at this iteration. A conservative selection approach takes a combination of small γ and T , whereas a more aggressive approach takes a combination of larger value γ and T . In the case when $\zeta = 1$, the MEBoost algorithm only updates the element(s) with the maximum absolute value. For any combination of γ and T , this is the most conservative approach that can be taken and will lead to sparser models than when a threshold is considered. It also requires a much larger value of T .

The parameter ζ determines how many coefficients are updated at each iteration; it offers a compromise between updating each coefficient at every iteration ($\zeta = 0$, similar to standard gradient descent) and updating only the coefficient corresponding to the element of the estimating equation with largest magnitude ($\zeta = 1$). In the context of Generalized Linear Models without measurement error, Wolfson (2011) showed that setting $\zeta = 1$ yields an update rule that is asymptotically equivalent (as $T \rightarrow \infty$, $\gamma \rightarrow 0$, and $T \cdot \gamma \rightarrow 0$) to following the path of minimizers of an L_1 -penalized projected artificial log-likelihood

ratio whose tangent is the GLM score function. By allowing multiple directions to be updated at each iteration, MEBoost can explore a much wider range of variable selection paths; as we now discuss later, cross-validation can be used to select the parameter ζ which leads to the optimal level of thresholding.

3.2.1 Connection to the ME-Lasso

Nakamura’s measurement corrected score functions are derivatives of corrected negative log likelihoods. In the normal case, the correction is exactly that described in Sorensen et al. (see Equation (2)). Hence, the arguments of ? can be applied to show that 1) MEBoost applied with S^* and threshold value $\tau = 1$, and 2) the solutions to (2), have the same local behavior. Specifically, under some regularity conditions (see Wolfson (2011)), as $T \rightarrow \infty$ and $\epsilon \rightarrow 0$ with $T \cdot \epsilon \rightarrow 0$, MEBoost’s iterative steps match the sequence of solutions to (2).

3.2.2 Selecting a final model

For a fixed ζ , identifying a final model involves choosing a point on the variable selection path generated by Algorithm 1; this is akin to choosing the penalty parameter in the Lasso. Cross-validation using a loss function relevant to the problem at hand (e.g., mean squared error) can be used to select a $\hat{\beta}$ on the path. Cross-validation can similarly be used to select the best value of ζ . The full procedure is described in Algorithm 2.

4 Simulation Study

To examine the impact of measurement error in the covariates on variable selection we performed a simulation study. We evaluated MEBoost by comparing it to two variable

Algorithm 2 Model Selection for MEBoost

procedure CROSS VALIDATION

Divide the observations into K folds where $\frac{1}{K}$ of the observations are used as a test set.

for $k = 1, \dots, K$ **do**

 Apply MEBoost for fixed value ζ .

 Obtain the mean squared prediction error of each candidate model on the test set.

 Calculate $\|\hat{\beta}\|_1^{(k)}$ for the $\hat{\beta}^{(t)}$ that minimizes mean squared prediction error.

Repeat across the K possible test sets and compute the mean of the selected $\|\hat{\beta}\|_1^{(k)}$'s.

selection methods: the Convex Conditioned Lasso (CoCoLasso), and the “naive” Lasso which does not correct for measurement error.

4.1 Simulation Set-up

Data were generated from a linear regression model with iid normal errors, $\mathbf{Y} = \mathbf{X}\beta + \epsilon$; where $\epsilon_1 \sim N(\mathbf{0}, \sigma_\epsilon^2)$ and $\sigma_\epsilon = 1.5$. The sample size for all studies is 80. The true covariates are drawn from a multivariate normal distribution, $\mathbf{X} \sim MVN(\mathbf{0}, \Sigma_{XX})$. Σ_{XX} is a block diagonal matrix with diagonal entries equal to 1, and 10 by 10 blocks corresponding to a group of 10 covariates with an exchangeable correlation structure with common pairwise correlation $\phi = 0.3$. In all simulations the true model has 10 non-zero coefficients and 90 zero coefficients, i.e., $\beta = (\mathbf{1}_{10}, \mathbf{0}_{90})$, so that the relevant covariates in the first block were correlated.

The measured covariates were generated as $\mathbf{W} = \mathbf{X} + \mathbf{U}$ for \mathbf{U} a matrix whose columns were generated as described below. To explore the impact of different types of measurement error, we considered 10 different scenarios for generating the columns of \mathbf{U} and varying the assumptions made about it. In the first five scenarios, \mathbf{U} is assumed to be normally distributed with mean zero and covariance matrix Ω , and the scenarios explore different structures for Ω . In each of Scenarios 1-5, we correctly specify the distribution

of \mathbf{U} when applying MEBoost and the CoCoLasso. Scenarios 6-10 explore cases where the distribution of \mathbf{U} is incorrectly specified.

1. Base case: $\mathbf{U} \sim N(\mathbf{0}, \delta^2 \Omega_1)$, where $\Omega_1 = \mathbf{I}$ the identity matrix, and $\delta^2 = 0.75$.
2. Varying δ^2 : $\delta_j^2 = 0.3375 + 0.075j$ for j in 1-10. This pattern repeats across the blocks of 10 covariates. The relevant covariates have similar spreads of measurement error to the irrelevant covariates.
3. Correlated Measurement Error: Measurement error was assumed normally distributed with an exchangeable correlation structure Within each block, $\mathbf{U} \sim N(\mathbf{0}, \delta^2 \Omega_3)$ where $\Omega_3 = \rho \mathbf{1}\mathbf{1}' + (1 - \rho)\mathbf{I}$, $\delta^2 = 0.75$, $\rho = 0.3$, and $\mathbf{1}$ is vector of ones.
4. Varying δ^2 with correlation: \mathbf{U} is distributed normally and is centered at 0 with $\delta_j^2 = 0.3375 + 0.075j$ for j in 1-10. This pattern repeats across the blocks of 10 covariates. The relevant covariates have similar spreads of measurement error to the irrelevant covariates. In blocks of ten, there is correlation in \mathbf{U} of 0.3
5. Some \mathbf{U} 's not measured with error: $\mathbf{U} \sim N(0, \delta^2 \Omega_5)$, where $\delta^2 = 0.75$, $\text{diag}(\Omega_5) = [0, 1, 0, 1, \dots]$ and $\Omega_{5,ij} = 0$ for $i \neq j$.
6. Overestimated δ^2 : \mathbf{U} generated as in Scenario 1, but we specify $\delta^2 = 1.5$.
7. Underestimated δ^2 : \mathbf{U} generated as in Scenario 1, but we specify $\delta^2 = 0.375$.
8. Misspecified correlation: \mathbf{U} generated as in Scenario 3, but we ignore the correlation and specify $\Omega = \delta^2 \mathbf{I}$ in running MEBoost and CoCoLasso.
9. Measurement error is distributed uniformly: Each entry \mathbf{U}_{ij} of \mathbf{U} is generated independently from a Uniform distribution, $\mathbf{U}_{ij} \sim U(-1.5, 1.5)$. MEBoost and CoCoLasso are run assuming $\mathbf{U} \sim N(\mathbf{0}, \delta^2 \mathbf{I})$ with $\delta^2 = 0.75 = \text{Var}(\mathbf{U}_{ij})$.

10. Measurement error is distributed asymmetrically: Each entry \mathbf{U}_{ij} of \mathbf{U} is generated independently from a shifted exponential distribution, $\mathbf{U}_{ij} + \sqrt{0.75} \sim \exp(\sqrt{0.75})$. MEBoost and CoCoLasso are run assuming $\mathbf{U} \sim N(\mathbf{0}, \delta^2 \mathbf{I})$ with $\delta^2 = 0.75 = \text{Var}(\mathbf{U}_{ij})$.

MEBoost was performed for each threshold value in the set $\zeta = (0.0, 0.2, 0.4, 0.6, 0.8, 1.0)$, and cross-validation (using the error-prone covariates) was used to select the optimal value of ζ and number of MEBoost iterations, as well as the value of λ in the CoCoLasso and naive Lasso. We compared MEBoost, CoCoLasso, and naive Lasso on two metrics of prediction error: mean squared error based on the true covariates ($\text{MSE} = \frac{1}{n}(\mathbf{Y} - \mathbf{X}\hat{\beta})'(\mathbf{Y} - \mathbf{X}\hat{\beta})$), mean squared error prediction based on the measured covariates ($\text{MSE-M} = \frac{1}{n}(\mathbf{Y} - \mathbf{W}\hat{\beta})'(\mathbf{Y} - \mathbf{W}\hat{\beta})$). These metrics were estimated using independent test sets generated during each individual simulation. We also computed L_1 distance from the true β , and variable selection sensitivity and specificity. For each scenario the metrics presented are the average over 1,000 simulations, and are calculated at intervals of 0.05 along $\|\hat{\beta}\|_1 \in \{0.05, 0.1, 0.15, \dots, 15\}$; the true value, $\|\beta\|_1 = 10$. Because the MEBoost algorithm may change multiple indices at each iteration it may not have values along each interval in the path. To account for this, a linear approximation of the relevant statistic was made at each point in the path.

We note that in this simulation study we chose to investigate model performance based on both the true and error-prone covariates. The motivation for techniques like ours which account for measurement error is to uncover the underlying relationship between the error-free covariates \mathbf{X} and the outcome \mathbf{Y} . Hence, in an ideal world, values of \mathbf{X} would be available on some subset (or an independent set) of observations so that prediction error could be assessed and the “best” model chosen. However, in practice

we will often only have access to the error-prone covariates \mathbf{W} for model fitting. So, if error-free measurements \mathbf{X} are not (and may never be) available, is it worthwhile to correct for measurement error? Buonaccorsi (1995) argued against correction, using the logic that the future predictions will be based on (error-prone) \mathbf{W} , not on (error-free) \mathbf{X} . Indeed, it can be shown in simple linear regression, that without the correction in a large sample the expected value of MSE-M is less than or equal to that of an estimate ignoring measurement error. However, as seen in the results section that follows, we found that correcting for measurement error decreased prediction error regardless of whether predictions were computed using error-free or error-prone covariates. This finding is discussed in greater detail below.

4.2 Simulation Results

Table 1 presents the minimum MSE, MSE-M, L_1 distance from the true β , sensitivity, and specificity at the minimum MSE for the three variable selection methods across the 10 scenarios. In all ten scenarios, MEBoost had the lowest MSE, MSE-M, and L_1 distance from the true β . The CoCoLasso has 16.6%-71.7% higher prediction error from the true covariates than MEBoost and in the case where measurement error is overestimated, the prediction error from the CoCoLasso is 5.26 times that of MEBoost. This is likely due to the fact that the Loh and Wainwright correction $\hat{\Sigma}$ in (5) is more negative, and hence requires a “longer” (and hence potentially more distorting) projection onto the space of positive definite matrices.

In terms of variable selection, MEBoost had a greater sensitivity and lower specificity than CoCoLasso in each case while Lasso had the lowest specificity. The Lasso struggles most when correlation is present in the measurement error. The MSE is about 2.5 times

that of MEBoost, when we allow MEBoost to account for the correlation. All methods perform poorly when we misspecify Δ by ignoring the correlation. The sensitivity and specificity are at high levels for most simulations with the exception of the misspecified Δ that ignored correlation. Overestimating δ lead to a more conservative selection process with a high specificity, while underestimating δ had a higher sensitivity. The L_1 distance from the true β can also tell us about performance. Again, the scenario where we misspecify Δ by ignoring correlation performs worst.

Figure 1 shows how these metrics vary with $\|\hat{\beta}\|_1$ in Scenario 2 (varying δ); similar plots for the other 9 scenarios appear in the supplementary materials. The MSE, MSE-M, and L_1 distance are U-shaped with a minimum just before $\|\hat{\beta}\|_1$ reaches the true $\|\beta\|_1$ of 10. For each metric, MEBoost achieves the lowest values, meaning it is closest to the true vector β and has the best prediction given either true covariates or covariates measured with error. The CoCoLasso performs worst. For sensitivity, MEBoost most quickly goes towards 1. This comes at the cost of specificity, but specificity does not begin to decrease drastically until after $\|\hat{\beta}\|_1 = 6$. The sensitivity of the Lasso eventually surpasses that of MEBoost, but when it does, the specificity decreases very sharply. The CoCoLasso at every point has the lowest value of sensitivity and the highest of specificity.

Figure 2 summarizes the differences between the variable selection path defined by MEBoost and CoCoLasso, in Scenario 1. We see in the left-hand plot that MEBoost includes nearly all relevant covariates within the first few iterations. The coefficient values are then increased at each step towards their true values of 1; at some point, additional irrelevant covariates begin to enter the model. In contrast, The CoCoLasso tends to increase the magnitude of a single coefficient, $\hat{\beta}_i$. This is confirmed by the middle plot, which shows the proportion of variables which are included in only one of

the two models, as $\|\hat{\beta}\|_1$ increases. We see from this plot that the number of covariates included in one model but not the other decreases as the CoCoLasso identifies more of the relevant covariates, then increases again as each method identifies different sets of irrelevant covariates. The rightmost plot shows the element-wise maximum and mean absolute difference between the estimated coefficient vectors from the two methods. We see that the CoCoLasso is quickly updating fewer elements as the maximum distance between estimates of β_i between the two methods grows initially, while MEBoost updates each of the relevant covariates gradually. As $\|\hat{\beta}\|_1$ approaches 4, the maximum distance begins to decrease. This is because MEBoost begins updating the index of the largest estimate from the CoCoLasso more quickly than the CoCoLasso is, since it has turned its attention to other variables. Once $\|\hat{\beta}\|_1$ surpasses the true value of 10, inclusion of irrelevant covariates causes the maximum distance to increase again.

5 Data Application

We applied our method to baseline data collected in the Box Lunch Study, a randomized trial of the effects of portion size availability on weight change. In the study, a total of 219 subjects were randomized to one of four groups: in three groups, subjects were provided a free daily lunch with a fixed number of calories (400, 800, and 1600). The control group was not provided a free lunch.

We considered the problem of predicting the number of times subjects reported bingeing on food in the last month, using Poisson regression with 99 explanatory variables. All variables were measured at baseline. 16 of the 99 explanatory variables were self-reported measures; of these 16, 8 were measures of food consumption and therefore possibly subject to substantial measurement error we will notate δ_D^2 . Another 8 may have also been

measured with error, notated δ_M^2 . Kipnis et al. (2003) examined a nutritional study with a 24 hour recall, and found that the correlation between the true and reported consumption of protein and energy was only 0.336. We assume this relationship exists in each of our variables measured with error. Assuming the measurement error variance $Var(U_i) \equiv \delta_i^2$ is independent of the variance of the true covariate $Var(X_i) \equiv \sigma_{X_i}^2$, we can obtain:

$$\rho_{W_i, X_i} = \rho_{X_i + U_i, X_i} = \frac{\sigma_{X_i}^2}{\sqrt{\sigma_{X_i}^2(\sigma_{X_i}^2 + \delta_i^2)}} \implies Var(W_i) = 1 - \rho_{X_i + U_i, X_i}^2 = \frac{\delta_i^2}{\sigma_{X_i}^2 + \delta_i^2} \quad (13)$$

and hence $Var(W_i) = 1 - 0.336^2 = 0.887$. This is the value we will need to provide MEBoost for our assumption of the measurement error. We assume this level of measurement error for each 24 hour dietary recall variable. After scaling our predictors to have zero mean and unit variance, we applied our method with the Nakamura correction. Since our measured data has its variance $(\delta_i^2 + \sigma_{X_i}^2)$ scaled to equal 1, we assumed that the 8 dietary recall covariates measured with error had $\hat{\delta}_D^2 = 0.887$. Since dietary variables may be more prone to measurement error than other variables, we scaled the assumed error of the other 8 variables to be half that of the nutritional variables: $\hat{\delta}_M^2 = \hat{\delta}_D^2/2$. The remaining variables were assumed to be measured without error. We conducted a sensitivity analysis to assess the performance of our method by setting $\hat{\delta}_D^2 = 0.5$ and 0.25.

To select tuning parameters, we employed 8-fold cross validation based on the deviance on a training set consisting of 70% of the data. The performance of our model was evaluated on the remaining test set. We present the models derived from MEBoost performed with three different thresholds ζ : 0.2, 0.6 (the approximate value estimated using cross-validation), and 0.9.

Table 2 shows the selected variables and estimated prediction error (MSE-M, bottom row) for various MEBoost models along with results from the naive Lasso. We did

not compare to the Measurement Error Lasso or the CoCoLasso because implementing these techniques in a problem of this size was computationally infeasible. The deviance and MSE-M were lowest for the model selected by MEBoost assuming the highest measurement error ($= 0.887$) and a threshold value of 0.6. This model ($\hat{\delta}_D^2 = 0.887$ and $\zeta = 0.6$) selected just 4 variables, which were a subset of the 7 chosen with the naive Lasso. The other two MEBoost models included up to two additional variables to the MEBoost model that minimized MSE-M (selected with $\hat{\delta}_D^2 = 0.887$ and $\zeta = 0.6$). Regardless of the assumption about the level of measurement error, using a threshold value of $\zeta = 0.2$ leads to the inclusion of several variables with small coefficients, and a much higher deviance and prediction error. Of particular note is that the naive Lasso (and MEBoost with the lower threshold) included the variable corresponding to the number of daily calories consumed at breakfast, while the best-performing MEBoost models (with $\zeta = 0.6$ and 0.9) did not. Since it is based on a 24-hour dietary recall, this variable may be particularly susceptible to measurement error induced by recall bias.

6 Discussion

We examined the variable selection problem in regression when the number of potential covariates is large compared to the sample size and when these potential covariates are measured with measurement error. We proposed MEBoost, a computationally simple descent-based approach which follows a path determined by measurement error-corrected estimating equations. We compared MEBoost, via simulation and in a real data example, with the recently-proposed Convex Conditioned Lasso (CoCoLasso) as well as the naive Lasso which assumes that covariates are measured without error. In almost all simulation scenarios, MEBoost performed best in terms of prediction error and coefficient bias.

The CoCoLasso is more conservative with the highest specificity in each case, but sensitivity and prediction are better with MEBoost. In the comparison of selection paths, we saw that MEBoost was more aggressive in identifying variables to be included in the model more quickly than the CoCoLasso. These differences were most apparent when the measurement error had a larger variance and a more complex correlation structure. In addition, when faced with a data set of 1000 observations and 1000 covariates, MEBoost obtained a solution in 1.3 seconds, while the CoCoLasso took 6:17.

As shown in the simulation study, MEBoost has lower prediction error than the Lasso on independent test data when predictions are based on the true (i.e., non-error-prone) covariates. It is interesting to note that MEBoost retains some advantage, albeit a more modest one, over the Lasso when predictions are based on error-prone covariates. This finding appears to contradict the intuition that accounting for covariate measurement error provides no benefit when the goal is prediction and error-free covariates will never be available. However, the observed benefit in our simulation is likely due to the fact that MEBoost is somewhat more flexible than the Lasso as it uses an additional parameter, the threshold ζ , which allows it to explore the model space more comprehensively.

MEBoost, while a promising approach, has some limitations. One limitation—which is shared with many methods that correct for measurement error—is that we assume that the covariance matrix of the measurement error process is known, an assumption which in many settings may be unrealistic. In some cases, it may be possible to estimate these structures using external data sources, but absent such data one could perform a sensitivity analysis with different measurement error variances and correlation structures, as we demonstrate in the real data application. Another challenging aspect of model selection with error-prone covariates is that, even if the set of candidate models is generated via

a technique which accounts for measurement error, the process of selecting a final model (e.g., via cross-validation) still uses covariates that are measured with error. However, we showed in our simulation study that MEBoost performs well in selecting a model which recovers the relationship between the true (error-free) covariates and the outcome, even when using error-prone covariates to select the final model. This finding suggests that the procedure for generating a “path” of candidate models has a greater influence on prediction error and variable selection accuracy than the procedure picking a final model from among those candidates.

To conclude, we note that while we only considered linear and Poisson regression in this paper, MEBoost can easily be applied to other regression models by, e.g., using the estimating equations presented by Nakamura (1990) or others which correct for measurement error. In contrast, the approaches of Sørensen et al. (2012) and Datta and Zou (2017) exploit the structure of the linear regression model and it is not obvious how they could be extended to the broader family of generalized linear models. The robustness and simplicity of MEBoost, along with its strong performance against other methods in the linear model case suggests that this novel method is a reliable way to deal with variable selection in the presence of measurement error.

References

- Brown, B., Miller, C., and Wolfson, J. (2016), “ThrEEBoost: Thresholded Boosting for Variable Selection and Prediction via Estimating Equations ,” *Journal of Computational and Graphical Statistics*.
- Buonaccorsi, J. (2010), *Measurement Error: Models, Methods and Applications*, Boca

Raton: CRC Press.

Buonaccorsi, J. P. (1995), “Prediction in the Presence of Measurement Error: General Discussion and an Example Predicting Defoliation,” *Source: Biometrics*, 51, 1562–1569.

Datta, A. and Zou, H. (2017), “CoCoLasso for High-dimensional Error-in-variables Regression,” *Annals of Statistics*.

Fraser, G. E. and Stram, D. O. (2012), “Regression calibration when foods (measured with error) are the variables of interest: markedly non-Gaussian data with many zeroes.” *American journal of epidemiology*, 175, 325–31.

Fuller, W. A. (ed.) (1987), *Measurement Error Models*, Wiley Series in Probability and Statistics, Hoboken, NJ, USA: John Wiley & Sons, Inc.

Kipnis, V., Subar, A. F., Midthune, D., Freedman, L. S., Ballard-Barbash, R., Troiano, R. P., Bingham, S., Schoeller, D. A., Schatzkin, A., and Carroll, R. J. (2003), “Structure of dietary measurement error: results of the OPEN biomarker study.” *American journal of epidemiology*, 158, 14–21; discussion 22–6.

Liang, K.-Y. and Zeger, S. L. (1986), “Longitudinal data analysis using generalized linear models,” *Biometrika*, 73, 13–22.

Loh, P.-L. and Wainwright, M. J. (2012), “High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity,” *The Annals of Statistics*, 40, 1637–1664.

- Nakamura, T. (1990), “Corrected score function for errors-in-variables models: Methodology and application to generalized linear models,” *Biometrika*, 77, 127–137.
- Rosner, B., Spiegelman, D., and Willett, W. C. (1992), “Correction of logistic regression relative risk estimates and confidence intervals for random within-person measurement error.” *American journal of epidemiology*, 136, 1400–13.
- Sørensen, Ø., Frigessi, A., and Thoresen, M. (2012), “Measurement Error in Lasso: Impact and Correction,” *arXiv.org*.
- Spiegelman, D., McDermott, A., and Rosner, B. (1997), “Regression calibration method for correcting measurement-error bias in nutritional epidemiology.” *The American journal of clinical nutrition*, 65, 1179S–1186S.
- Stefanski, L. A. and Carroll, R. J. (1985), “Covariate Measurement Error in Logistic Regression,” *The Annals of Statistics*, 13, 1335–1351.
- Tibshirani, R. (1996), “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society Series B*, 58, 267–288.
- Wolfson, J. (2011), “EEBoost: A General Method for Prediction and Variable Selection Based on Estimating Equations,” *Journal of the American Statistical Association*, 106, 296–305.

7 Tables and Figures

Scenario	Method	MSE	MSE-M	L_1D	SENS	SPEC
Measurement error iid	MEBoost	4.86	10.65	5.17	0.95	0.86
	Lasso	7.13	11.63	6.75	0.98	0.75
	CoCoLasso	6.30	12.53	6.04	0.92	0.91
Varying δ	MEBoost	4.88	10.52	5.21	0.96	0.85
	Lasso	7.08	11.42	6.76	0.98	0.76
	CoCoLasso	7.18	14.84	6.50	0.85	0.95
Some $\delta = 0$	MEBoost	3.65	6.70	3.57	0.99	0.87
	Lasso	4.88	6.23	5.42	0.99	0.83
	CoCoLasso	6.23	11.15	5.60	0.92	0.95
Varying δ & correlation	MEBoost	6.19	19.12	6.27	0.95	0.87
	Lasso	15.18	21.42	9.63	0.80	0.79
	CoCoLasso	8.94	22.47	7.77	0.78	0.94
Correlation in measurement error	MEBoost	6.16	19.27	6.29	0.95	0.87
	Lasso	15.78	22.24	9.75	0.80	0.79
	CoCoLasso	8.67	22.29	7.81	0.78	0.94
Overestimated δ	MEBoost	4.00	10.21	3.46	0.94	0.94
	Lasso	7.18	11.71	6.75	0.98	0.76
	CoCoLasso	21.05	28.22	9.11	0.37	1.00
Underestimated δ	MEBoost	5.54	10.81	6.28	0.98	0.76
	Lasso	7.18	11.71	6.75	0.98	0.76
	CoCoLasso	6.46	12.02	6.20	0.95	0.87
Misspecified Δ , ignores correlation	MEBoost	12.79	21.03	9.41	0.86	0.80
	Lasso	15.78	22.24	9.75	0.80	0.79
	CoCoLasso	15.67	24.17	9.51	0.55	0.93
Measurement error from asymmetric distribution	MEBoost	4.89	10.68	5.16	0.95	0.85
	Lasso	7.21	11.85	6.81	0.98	0.75
	CoCoLasso	7.32	15.20	6.71	0.85	0.95
Measurement error from uniform distribution	MEBoost	4.81	10.52	5.17	0.96	0.84
	Lasso	7.19	11.75	6.80	0.99	0.75
	CoCoLasso	6.61	13.91	6.26	0.87	0.94

Table 1: Performance metrics for the 1,000 simulations in various measurement error scenarios. The models were selected at the point with minimum MSE-M.

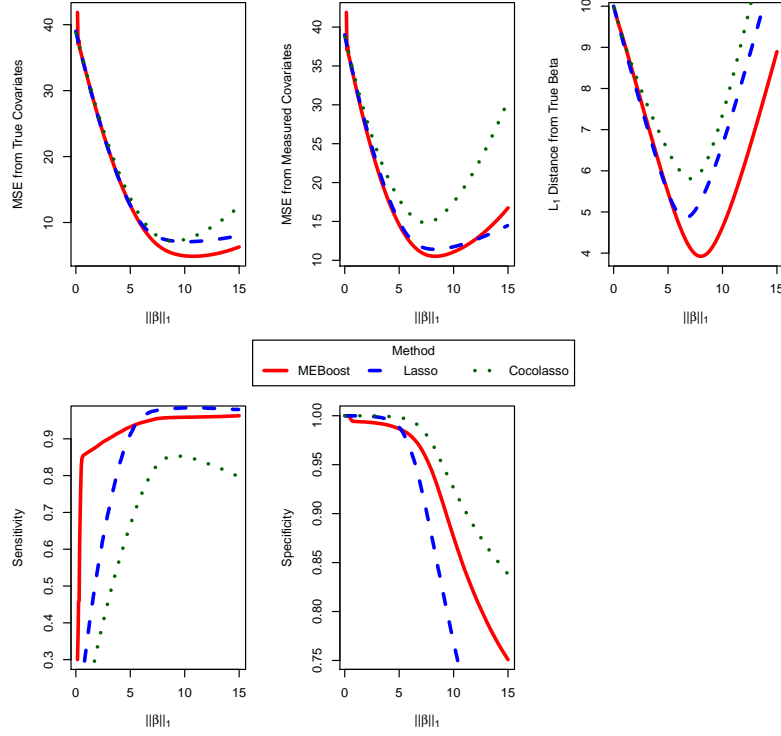


Figure 1: Summary statistics for the scenario with varying levels of independent measurement errors. Plots are of MSE, MSE-M, L_1 distance from $\hat{\beta}$ to the true value of β , sensitivity, and specificity across the mean path over 1,000 simulations.




Figure 2: In the scenario, with iid measurement error, 1,000 simulations were conducted comparing the variable selection path of MEBoost to CoCoLasso. Plots are of the number of nonzero coefficients included in the model, the portion of coefficients that are 'mismatched' with only one of the two models selecting a certain variable, and the element-wise distance in a coefficient between the two methods.

Variable	$\hat{\delta}_D^2 = 0.887$			$\hat{\delta}_D^2 = 0.5$			$\hat{\delta}_D^2 = 0.25$			Naive	
	$\zeta = 0.2$	$\zeta = 0.6$	$\zeta = 0.9$	$\zeta = 0.2$	$\zeta = 0.6$	$\zeta = 0.9$	$\zeta = 0.2$	$\zeta = 0.6$	$\zeta = 0.9$	Lasso	
Ate lg amt past 28 days	0.09	0.23	0.39	0.08	0.23	0.33	0.09	0.22	0.29	0.21	
Lost control past 28 days	0.09	0.23	0.21	0.08	0.23	0.22	0.09	0.22	0.23	0.24	
TFEQ Disinhibition	0.09	0.23	0.09	0.08	0.23	0.18	0.09	0.22	0.21	0.17	
BCT: Max clicks for pizza slice	0.09	0.21	0.1	0.08	0.23	0.19	0.09	0.22	0.2	0.16	
Long fast (0=no days 6=every day)	0.09	-	-	0.08	-	-	0.09	-	-	-	
Judge your shape (0=not at all 6=markedly)	0.07	-	-	0.07	0.06	-	0.07	0.06	-	-	
Judge your weight (0=not at all 6=markedly)	0.07	-	-	0.07	-	0.06	0.07	0.06	0.06	0.06	
Disatisfied with shape (0=not at all 6=markedly)	0.07	-	-	0.07	-	-	0.07	-	-	-	
BCT: Rounds of clicking for pizza	0.06	-	-	0.06	-	-	0.07	-	-	-	
BCT: Pmaxpizza/(Pmaxpizza+Pmaxread)	0.06	-	-	0.06	-	-	0.06	-	-	-	
TFEQ Hunger	0.06	-	-	0.06	-	-	0.06	-	-	-	
Dissatisfied with weight (0=not at all 6=markedly)	0.06	-	-	0.06	-	-	0.06	-	-	-	
NDSR breakfast kcals at BL	0.06	-	-	0.05	-	-	-	-	-	0.08	
CDRS body image (1=thinnest 9=fattest)	0.05	-	-	0.05	-	-	0.05	-	-	-	
Dem9 Household income	-	-	-	-	-	-	-	-	-	-0.05	
Eat lunch in cafeteria	-	-	-	-	-	-	-0.05	-	-	-	
Eat items from home, days/wk	-0.06	-	-	-0.06	-	-	-0.06	-	-	-	
Cohort first lunch date	-0.07	-	-	-0.07	-	-	-0.08	-	-	-	
Deviance	210.43	89.06	100.49	222.07	90.22	93.23	208.07	90.63	90.38	97.12	
MSE-M	22.35	5.61	8.10	23.28	5.69	7.29	22.04	5.97	6.21	7.89	

Table 2: Coefficients, Deviance, and MSE-M from selected models for MEBoost with specified value of ζ and $\hat{\delta}_D^2$ and the Lasso. Small coefficients (magnitude < 0.05) are omitted. “-” indicates that the variable was not selected in the model.